arXiv:1601.03768v1 [cs.SY] 14 Jan 2016

Optimal Current Waveforms for Switched-Reluctance Motors

Nicholas Moehle* Stephen Boyd†

January 18, 2016

Abstract

In this paper, we address the problem of finding current waveforms for a switched reluctance motor that minimize a user-defined combination of torque ripple and RMS current. The motor model we use is fairly general, and includes magnetic saturation, voltage and current limits, and highly coupled magnetics (and therefore, unconventional geometries and winding patterns). We solve this problem by approximating it as a mixed-integer convex program, which we solve globally using branch and bound. We demonstrate our approach on an experimentally verified model of a fully pitched switched reluctance motor, for which we find the globally optimal waveforms, even for high rotor speeds.

1 Introduction

We consider the problem of choosing optimal current waveforms for a switched reluctance motor (SRM). Traditionally, the shape of the current waveforms is determined in an ad-hoc manner (*e.g.*, by fixing the turn-on angles for each phase winding as a function of rotor position; see [TK12]). Although creating these waveforms does not require a detailed motor model, and implementing them is simple, waveforms produced in this manner rarely produce smooth output torque, and often decrease motor efficiency and exacerbate mechanical vibration and acoustic noise issues. Also, because of inverter voltage limits, such waveforms may not even be realizable at high rotor speeds.

We therefore propose to use optimization to find current waveforms that acheive a desired average torque, while minimizing a combination of resistive power loss and RMS torque ripple. We consider supply voltage limits, as well as current limits in each phase winding. Our model includes a detailed magnetic circuit, which can account for magnetic coupling between phases, and can be used to model motors with unconventional geometries and winding patterns, such as those with fully pitched windings.

We propose to solve this optimization problem by approximating it as a mixed-integer convex program (MICP). This MICP reformulation approach has several inherent advantages over more conventional methods, such as sequential quadratic programming. The most prominent, for our purposes, is that it can be solved globally by generic methods such as branch and bound, often in a reasonable amount of time. This has two benefits: first, it allows us to achieve the best performance possible for a given motor; second, it provides a metric against which other, suboptimal methods

*Mechanical Engineering Department, Stanford University. moehle@stanford.edu
†Electrical Engineering Department, Stanford University. boyd@stanford.edu

can be compared. We note that although in general global optimization methods for solving MICPs can have very high (exponential) runtime, we find that global solutions can typically be found in a reasonable amount of time (1-5 minutes) for the problems we encounter. These waveforms can be computed and stored in a lookup table, indexed by desired torque and rotor speed. Such a table with hundreds or thousands of entries could be computed overnight on a standard multi-core computer.

Another advantage of an MICP formulation is that, if suboptimal solutions are acceptable, first-order methods exist that can produce a good solution very quickly especially when initialized with a decent initial guess. (see, *i.e.*, [TMBB15]) This opens the possibility that (nearly) optimal waveforms can be produced by a smart control scheme even as motor parameters vary over the life of the motor. This is especially attractive when combined with a lookup table containing precomputed, globally optimal waveforms: the precomputed waveforms can be used as an initial guess for an online optimization method when new waveforms (corresponding to updated motor parameters) are required.

We demonstrate our approach numerically on the experimentally validated motor model of [MWC01], which describes an SRM with fully pitched phases.

1.1 Previous work

Current optimization for SRMs. Several authors have considered optimization of SRM current waveforms. The most similar work to our own is a series of papers by Lovatt and Stephenson [LS94], [LS97a], [LS97b], that use local optimization methods to find minimum RMS current waveforms that acheive a given (average or pointwise) desired torque, subject to voltage and current constraints. Stankovic et al. derive optimal waveforms for a simple SRM, under several restrictive assumptions (*e.g.*, sufficient drive voltage, no more than two simultaneously conducting phases); under these assumptions, it is only necessary to discretize the waveforms during commutation. Kaiserseder et al. [KSAS03] also seek optimal current waveforms that produce smooth torque output and can be chosen to either minimize current RMS values, or minimize vibration resulting from radial forces. The method of optimization is not described. A similar optimization problem is posed by Chapman and Sudhoff [CS02] in the frequency domain; sequential quadratic programming is used to (approximately) solve it. Many other works optimize over current waveforms parametrized by only a few free variables, such as the firing angle or the corner locations of a trapezoidal waveform. This of course requires predetermined current waveform shapes, which are not optimal in general. For some examples of this approach, see [MK03], [CMH93], and [CKKP02].

We also note that our approach can be viewed as an extension of the authors' previous work on optimal current waveform design for permanent magnet motors; see [MB15].

Hybrid control and MPC for SRMs. Here we list some other optimization-based techniques applied to control SRMs. Peyrl et al. propose a finite-set model predictive control approach, which involves online optimization directly over future inverter switching states. Due to the computationally demanding nature of this technique, only very short prediction horizons (*i.e.*, up to three steps) are considered. The work by Vasak et al. [VZP+07], uses a piecewise-affine model of the torque characteristic (as we do) to derive a feedback controller that guarantees a torque ripple. However, their proposed model is relatively low fidelity: the dynamics are modelled as a first-order linear system, and the proposed torque characteristic (our $g_k(\mathcal{F}_k, \theta)$) has only nine regions (For comparison, in §6 we use over one hundred regions.)

MICP. Convex optimization problems can be solved efficiently and reliably using standard techniques [BV04] (and additionally, specialized modelling software, such as CVX [GB14], enables rapid development of convex optimization applications). An optimization problem with some integer variables, but which is otherwise convex, is called a mixed integer convex program (MICP). Because of the presence of integer variables, MICPs are nonconvex optimization problems, and are difficult to solve (globally) in general (*i.e.*, these problems are are NP-hard; see [KT06]). Indeed, all known global solution techniques for MICPs (such as the branch-and-bound algorithm), have exponential worst-case runtime. Nevertheless, many of these algorithms are effective in practice, and we found them to work well for the the optimization problem we formulate in this paper. In addition to global solution algorithms, many approaches exist to (approximately) solve MICPs more quickly; see [TMBB15] and references therein.

Our formulation is based on approximating the nonlinear constraint functions (the torque characteristics and the magnetic flux characteristics) by piecewise affine functions. These constraints can then be represented as a combination of integer and linear constraints using disjunctive programming. For details on disjunctive programming consult Balas [Bal79] and Ceria and Soares [CS99]. Disjunctive programming has found many applications in the past decade or so, such as process engineering [GT13], facility location, unit commitment and portfolio management [GL12], and optimal control [**?**].

1.2 Contribution

Our reformulation of the torque control problem as a MICP opens the door for two interesting possiblilites. The first is that global solution methods can be used to find the optimal waveforms, typically in a reasonable amount of time. This is of course advantageous in its own right, as it allows provably optimal waveforms to be implemented. It is also useful to verify the limits of performance of motors, and as a benchmark for comparing heuristic methods. The second possibility is that fast first-order methods, which are simple enough to run on embedded platforms, can be used as a heuristic to update the globally optimal waveforms as parameters vary over the life of the motor. Our proposed model is also much more general than the optimization models considered in previous works, and can therefore be used to capture more of the characteristc features of switched reluctance motors, such as magnetic coupling. We hope this generality will be useful for researchers investigating novel switched reluctance motor topologies, by giving them a practical method for optimal waveform generation, and by characterizing the theoretical performance of their designs.

2 Motor model

We consider an abstract, lumped parameter model of a switched reluctance motor. The rotor, which does not contain any windings or magnetic elements, has angular position θ and angular velocity ω; we assume ω is constant. The stator contains n electrical circuit branches, called *windings*. The winding currents are $i \in \mathbf{R}^n$, the winding voltages are $v \in \mathbf{R}^n$, and the magnetic flux linkages through the windings are $\lambda \in \mathbf{R}^n$. The stator also contains m magnetic elements, with magnetomotive force (MMF) vector $\mathcal{F} \in \mathbf{R}^m$, and magnetic flux vector $\psi \in \mathbf{R}^m$.

We will assume that i, v, λ, $\mathcal{F}$, and ψ are 2π-periodic functions of θ. We use a prime ($'$) to denote differentiation of these functions with respect to θ. To lighten notation, we often drop explicit dependence on θ.

Electrical dynamics. The electrical circuit dynamics are

$$v = Ri + \omega\lambda', \tag{1}$$

where $R \in \mathbf{S}^n_{++}$ is the (diagonal) resistance matrix. Note that $\omega\lambda'$ is the time derivative of λ.

Magnetic circuit. We assume the magnetic elements are connected by a (planar) magnetic circuit, which we describe in terms of mesh analysis (for an introduction to mesh analysis, see [DK84]). In particular, we assume that there are l circuit meshes (not including the outer mesh) each with an associated reference direction. The (reduced) *mesh matrix* $M \in \mathbf{R}^{l\times m}$ is such that

$$M_{jk} = \begin{cases} 1 & \text{if magnetic element } k \text{ is in mesh } j, \text{ with coinciding reference directions} \\ -1 & \text{if magnetic element } k \text{ is in mesh } j, \text{ with opposite reference directions} \\ 0 & \text{if magnetic element } k \text{ is not in mesh } j. \end{cases}$$

Any flux vector ψ consistent with the magnetic circuit topology must be a linear combination of the rows of M, so that

$$\psi = M^T\phi, \tag{2}$$

for some $\phi(\theta) \in \mathbf{R}^l$, which we call the *mesh magnetic flux vector*. (Because M has full row rank in general, there is a unique ϕ for any such ψ.)

In addition, the total MMF around each mesh is the sum of the MMFs of the magnetic elements that make up the mesh (accounting for reference direction), so the vector of mesh MMFs is given by the vector $M\mathcal{F}$.

Electro-magnetic geometry. We define the *electro-magnetic geometry matrix* $C \in \mathbf{R}^{l\times n}$ such that C_{jk} gives the amount of current passing through mesh j per unit of current in winding k. The total current passing through each of the l meshes is therefore given by the vector Ci, which is related to the total MMF around the meshes by Ampère's law:

$$M\mathcal{F} = Ci. \tag{3}$$

Similarly, the flux linkage is related to the mesh magnetic flux vector by

$$\lambda = C^T\phi. \tag{4}$$

Magnetic characteristic. The magnetic flux and the MMF of the k-th magnetic element are related by

$$\psi_k(\theta) = f_k\big(\mathcal{F}_k(\theta), \theta\big), \tag{5}$$

where the magnetic characteristic f_k is a monotonically increasing function in its first argument.

As a special case, if the functions f_k are affine in $\mathcal{F}_k$ for each θ, with constant linear term (so that $\psi(\theta) = A\mathcal{F}(\theta) + b(\theta)$, with A diagonal and positive definite), as in the case of a permanent magnet motor, then we have $\lambda(\theta) = Li(\theta) + k(\theta)$, where $L = C^T(MA^{-1}M^T)^{-1}C$ is the *inductance matrix* and $k(\theta) = C^T(MA^{-1}M^T)^{-1}A^{-1}Mb(\theta)$ is the *back-emf constant*.

Torque. The magnetic co-energy is

$$E^*(\mathcal{F},\theta) = \sum_{k=1}^{m} \int_0^{\mathcal{F}_k} f_k(x,\theta)\, dx.$$

The electromagnetic torque can be expressed as

$$\tau(\theta) = -\frac{\partial}{\partial\theta} E^*\big(\mathcal{F}(\theta),\theta\big) = -\sum_{k=1}^{m} \frac{\partial}{\partial\theta} \int_0^{\mathcal{F}_k} f_k(x,\theta)\, dx.$$

By introducing a *phase torque* function g_k such that

$$g_k(y,\theta) = -\frac{\partial}{\partial\theta} \int_0^{y} f_k(x,\theta)\, dx,$$

we have

$$\tau(\theta) = \sum_{k=1}^{m} g_k\big(\mathcal{F}_k(\theta),\theta\big). \tag{6}$$

Voltage limits. We assume the winding voltages are bounded:

$$|v_k(\theta)| \le v^{\max}, \quad k = 1,\dots,n. \tag{7}$$

Torque ripple. The average torque over one cycle is

$$\overline{\tau} = \frac{1}{2\pi} \int_0^{2\pi} \tau(\theta)\, d\theta.$$

The (quadratic) torque ripple is

$$r = \frac{1}{2\pi} \int_0^{2\pi} \big(\tau(\theta) - \overline{\tau}\big)^2\, d\theta.$$

Power loss. The power loss is the average resistive loss from the phase currents over one cycle:

$$P_{\text{loss}} = \frac{1}{2\pi} \int_0^{2\pi} i(\theta)^T R i(\theta)\, d\theta.$$

3 Optimal torque control

The optimal torque control problem is to choose the phase voltages, phase currents, and eddy currents to achieve a desired average torque while minimizing the average power loss and torque ripple:

$$\begin{array}{ll} \text{minimize} & P_{\text{loss}} + \alpha r \\ \text{subject to} & \overline{\tau} = \tau^{\text{des}}, \\ & \text{equations (1), (2), (3), (4), (5), (6), and (7).} \end{array} \tag{8}$$

The parameters are the trade-off parameter $\alpha \ge 0$, the rotor angular velocity ω, the desired average torque τ^{des}, the resistance matrix R, the mesh matrix M, the electro-magnetic geometry matrix

C, and the magnetic characteristic functions f_k, for $k = 1, \ldots, m$. The problem variables are the 2π-periodic functions i, v, λ, $\mathcal{F}$, ψ, and ϕ.

Problem (8) is an infinite-dimensional optimization problem. The problem is nonconvex due to the magnetic characteristic (5) and the torque relation (6), and therefore is expected to be difficult to solve globally.

One approach is to find a locally optimal solution, using common local optimization methods such as sequential quadratic programming Software that implements these methods is readily available; see [NW06].

In this paper, we pursue a different approach, and will instead show how to approach (8) by discretizing the variables, and (approximately) reformulating the problem as a mixed-integer convex program (MICP), a problem class for which efficient algorithms are available to find a good (or even globally optimal) solution.

We note that if the magnetic characteristic is affine, with only the offset depending on rotor position, the torque control problem reduces to a version of the formulation given in [MB15]. Our current problem can therefore be interpreted as an extension of that formulation to cover magnetic nonlinearities and reluctance torque.

4 Conversion to MICP

In this section we show how to convert (8) to an (infinite-dimensional) mixed-integer convex program. To do this, we use piecewise-affine approximations of the nonlinear equality constraints.

4.1 Approximation of magnetic characteristic

Here we approximate the equation magnetic characteristic (5) by a set of linear and integer constraints. We replace the constraint $\psi_k = f_k(\mathcal{F}_k, \theta)$ with the constraint

$$\psi_k = \tilde{f}_k(\mathcal{F}_k, \theta), \tag{9}$$

where $\tilde{f}_k$ is a piecewise affine approximation of f_k. In particular, we have

$$\tilde{f}_k(x, \theta) = \begin{cases} a_k^1(\theta)x + b_k^1(\theta) & \tilde{\mathcal{F}}_k^0 \leq x \leq \tilde{\mathcal{F}}_k^1 \\ \vdots & \vdots \\ a_k^N(\theta)x + b_k^N(\theta) & \tilde{\mathcal{F}}_k^{N-1} \leq x \leq \tilde{\mathcal{F}}_k^N, \end{cases}$$

where $a_k^1(\theta), \ldots, a_k^N(\theta)$ and $b_k^1(\theta), \ldots, b_k^N(\theta)$ parametrize the affine functions, and $\tilde{\mathcal{F}}_k^0, \ldots, \tilde{\mathcal{F}}_k^N$ are the boundaries of the affine regions, so that we have

$$f_k(x, \theta) \approx a_k^j(\theta)x + b_k^j(\theta)$$

if $\tilde{\mathcal{F}}_k^{j-1} \leq x \leq \tilde{\mathcal{F}}_k^j$.

By introducing additional variables $z_k^j(\theta)$ and $s_k^j(\theta)$, for $j = 1, \ldots, N$, and $k = 1, \ldots, m$, the approximate magnetic characteristic constraint (9) can be written as

$$\begin{gathered} \psi_k(\theta) = \sum_{j=1}^N a_k^j(\theta) z_k^j(\theta) + b_k^j(\theta) s_k^j(\theta) \qquad \mathcal{F}_k(\theta) = \sum_{j=1}^N z_k^j(\theta) \qquad \sum_{j=1}^N s_k^j = 1 \\ \tilde{\mathcal{F}}_k^{j-1} s_k^j(\theta) \leq z_k^j(\theta) \leq \tilde{\mathcal{F}}_k^j s_k^j(\theta) \qquad s_k^j(\theta) \in \{0, 1\}. \end{gathered} \tag{10}$$

4.2 Approximation of torque function

In the same way, we can approximate the torque constraint (6) using a set of linear and integer constraints. To do this, we first approximate each torque function g_k, for $k = 1, \ldots, m$, as a piecewise affine function $\tilde{g}_k$:

$$\tilde{g}_k(x, \theta) = \begin{cases} c_k^1(\theta)x + d_k^1(\theta) & \tilde{\mathcal{F}}_k^0 \leq x \leq \tilde{\mathcal{F}}_k^1 \\ \vdots & \vdots \\ c_k^N(\theta)x + d_k^N(\theta) & \tilde{\mathcal{F}}_k^{N-1} \leq x \leq \tilde{\mathcal{F}}_k^N, \end{cases}$$

where $c_k^1(\theta), \ldots, c_k^N(\theta)$ and $d_k^1(\theta), \ldots, d_k^N(\theta)$ parametrize the affine functions. Then the approximate torque constraint

$$\tau(\theta) = \sum_{k=1}^{m} \tilde{g}_k(\mathcal{F}_k, \theta)$$

can be included by appending

$$\tau(\theta) = \sum_{k=1}^{m} \sum_{j=1}^{N} c_k^j(\theta) z_k^j(\theta) + d_k^j(\theta) s_k^j(\theta) \tag{11}$$

to the constraints (10) above.

4.3 Improving the MICP formulation

By converting the nonlinear constraints to linear and integer constraints, we have succeeded in our goal of making (8) into an MICP. However, the runtime of a global MICP solver often depends crucially on the problem formulation. Here we give an additional reformulation of the objective of (8) (specifically, of the the power loss) that may improve the performance of a MICP solver compared with the basic formulation.

Assuming that C has full column rank, we can use (3) to express the winding current as $i = C^\dagger M \mathcal{F}$, where $C^\dagger$ is the pseudo-inverse of C (or any other left inverse). Then the power loss can be rewriten as

$$P_{\text{loss}} = \frac{1}{2\pi} \int_0^{2\pi} \mathcal{F}(\theta)^T Q \mathcal{F}(\theta) \, d\theta,$$

where $Q = M^T (C^\dagger)^T R C^\dagger M$.

For any feasible set of variables, and for any diagonal matrix D, this is equivalent to

$$\frac{1}{2\pi} \int_0^{2\pi} \left(\mathcal{F}(\theta)^T (Q - D) \mathcal{F}(\theta) + \sum_{k=1}^{m} \sum_{j=1}^{N} D_{kk} \frac{z_k^j(\theta)^2}{s_k^j(\theta)} \right) d\theta.$$

In particular, if we choose D with nonnegative diagonal elements, such that $Q - D$ is positive semidefinite, then the reformulated power loss function is convex in all its variables. Furthermore, the larger the elements of D are, the tighter the convex relaxation will be. Finding a suitable D can be cast as a small convex optimization problem; for details on this type of reformulation, see [FG07].

5 Symmetry and discretization

5.1 Symmetry

Many motors have substantial symmetry, which in our formulation is encoded in the functions f_k, as well as M, R, C. Although it is *not* necessarily true that optimal variables for (8) share this symmetry, it is reasonable to expect that many good sets of variables are symmetric. Furthermore, we may explicitly desire that the waveforms be symmetric (*e.g.*, to simplify implementation, or to wear components evenly). By introducing symmetry constraints, we can also reduce the interval of the variables of (8) thus reducing the complexity of a discretized version of the problem. For some asymmetric motors, one or more of these assumptions may not hold; examples of this include motors intentionally designed without symmetry, or when a winding in an otherwise symmetric motor has failed.

Pole symmetry. We assume the rotor has N_p pole pairs *i.e.*, f_k is $2\pi/N_\mathrm{p}$-periodic for all k. Consequently, we restrict our search to variables i, v, λ, ψ, ϕ, and $\mathcal{F}$ that are also $2\pi/N_\mathrm{p}$-periodic.

Phase symmetry. We assume the motor has K phases, and therefore search for variables that satisfy the periodicity property

$$i_k(\theta) = i_1\left(\theta + \frac{2\pi(k-1)}{KN_\mathrm{p}}\right),$$

with similary constraints holding for v, λ, $\mathcal{F}$, ψ, and ϕ.

Equivalent problem. The symmetry assumptions allow us to form an equivalent problem with the same constraints and objective as (8) in which the variables have domain $[0, 2\pi/(KN_\mathrm{p})]$. We also add periodicity constraints of the form

$$i_k(0) = i_1\left(\frac{2\pi(k-1)}{KN_\mathrm{p}}\right), \qquad \text{for } k = 1, \ldots K.$$

with similar constraints for the other variables. The integrands in the definitions of average torque, torque ripple, and power loss are each $\pi/(3N_\mathrm{p})$-periodic; to get equivalent definitions of these values over the appropriate domain, we can integrate over $[0, 2\pi/(KN_\mathrm{p})]$ instead of $[0, 2\pi]$, and scale the result by KN_p.

5.2 Discretization

After reducing the domain of the variables of (8), we discretize this interval into $T+1$ grid points, $\theta_0, \ldots, \theta_T$, with $\theta_0 = 0$ and $\theta_T = 2\pi/(KN_\mathrm{p})$. All pointwise constraints must hold at $\theta_0, \ldots, \theta_{T-1}$, and the periodicity constraints must hold at θ_0 and θ_T. Integration over the interval is replaced by summation from θ_0 to θ_{T-1}, with appropriate scaling. We approximate the derivative in (1) using a forward difference approximation.

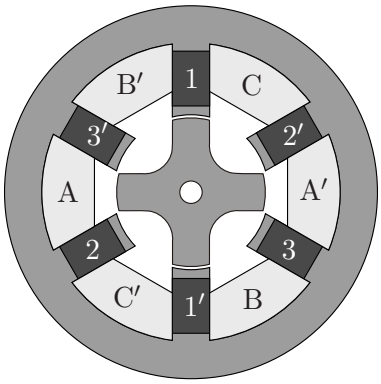

Figure 1: A switched reluctance motor with fully pitched windings. Phase windings are shown in light gray, and are labeled A, B, and C for current flowing out of the page (with complementary flow denoted with a prime). Magnetic elements are shown in dark gray, and are labeled 1, 2, and 3, with magnetic flow toward the rotor (with complementary flow denoted with a prime). Note that every third stator teeth belongs to a single magnetic element.

6 Example

6.1 Motor model

We consider the switched reluctance motor with fully pitched windings shown in figure 1. Our model is based the first example of [MWC01]. The motor model has fully pitched windings, instead of the more usual concentrated windings. (This is reflected by the matrix M, which would be diagonal for a motor with concentrated windings).

Symmetry and discretization. We assume that each stator tooth is magnetically identical, which allows us to group every third stator tooth into a single magnetic element. With this assumption, the motor exhibits pole symmetry, with $N_{\mathrm{p}} = 2$, and phase symmetry, with $K = 3$ phases. We discretize the interval $[0, \pi/3]$ with $T = 40$, so that $\theta_0 = 0$ and $\theta_T = \pi/3$.

Electromagnetic. We have $m = 3$ magnetic elements, and $l = 3$ meshes. The mesh matrix of the magnetic circuit is

$$M = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{bmatrix}.$$

The motor has $n = 3$ windings. The voltage limit is 600 V. The resistance and electro-magnetic geometry matrices are

$$R = \begin{bmatrix} 0.1 & 0 & 0 \\ 0 & 0.1 & 0 \\ 0 & 0 & 0.1 \end{bmatrix} \Omega, \qquad C = (1/2)N_{\mathrm{turns}} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

where $N_{\mathrm{turns}} = 204$ is the number of turns in each winding.

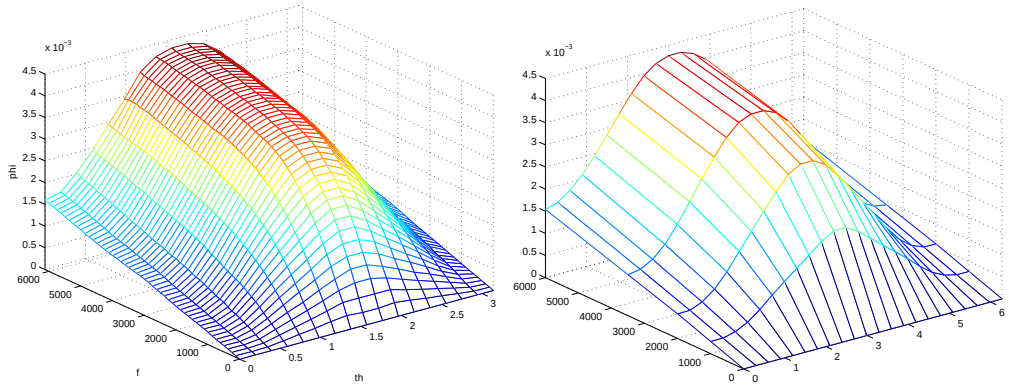

Figure 2: The flux characteristic f_1 (left), and its piecewise affine approximation $\tilde{f}_1$ (right).

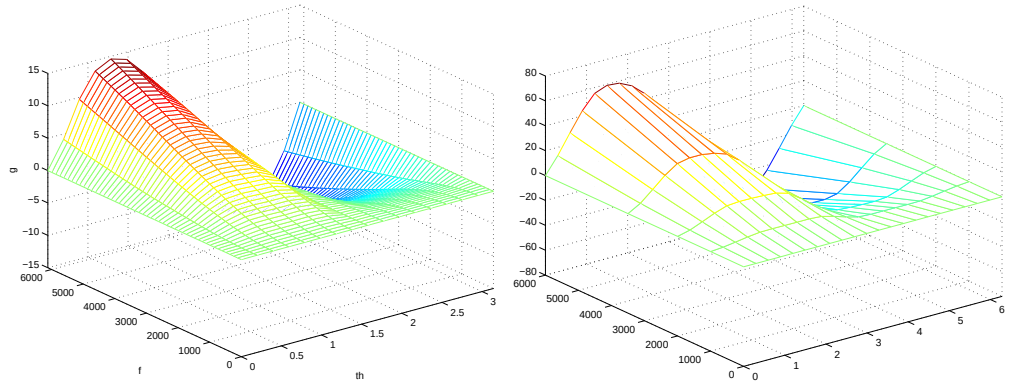

Figure 3: The torque characteristic g_1 (left), and its piecewise-affine approximation $\tilde{g}_1$ (right).

Flux characteristic and phase torque functions. The magnetic flux characteristic f_1 and the phase torque function g_1 are shown in figures 2 and 3. We only show these functions for the first magnetic element. (Due to phase symmetry, the other two functions are shifted versions of the first.)

For each value of θ_t, for $t = 0, \ldots, T$, we fit piecewise-affine functions to f_k and g_k. The piecewise-affine region boundaries were chosen as

$$\tilde{\mathcal{F}} = (0, 1000, 2000, 3500, 6000) \text{ Ampere-turns.}$$

With the region boundaries fixed, the values of the coefficient functions a_k, b_k, c_k, and d_k were determined at the relevant (discrete) rotor positions using least-squares, with the constraint that both piecewise-affine functions be continuous in $\mathcal{F}$ for all values of θ.

6.2 Results

Low-speed operation. The optimal current, voltage, and phase torque waveforms for $\omega = 1000$ rpm, $\tau_{\text{des}} = 10$ N $\cdot$ m and $\alpha = 3$ J/(N $\cdot$ m)^2 are shown in figure 4. (We restricted our search to waveforms that are phase symmetric, as discussed in §5.1.)

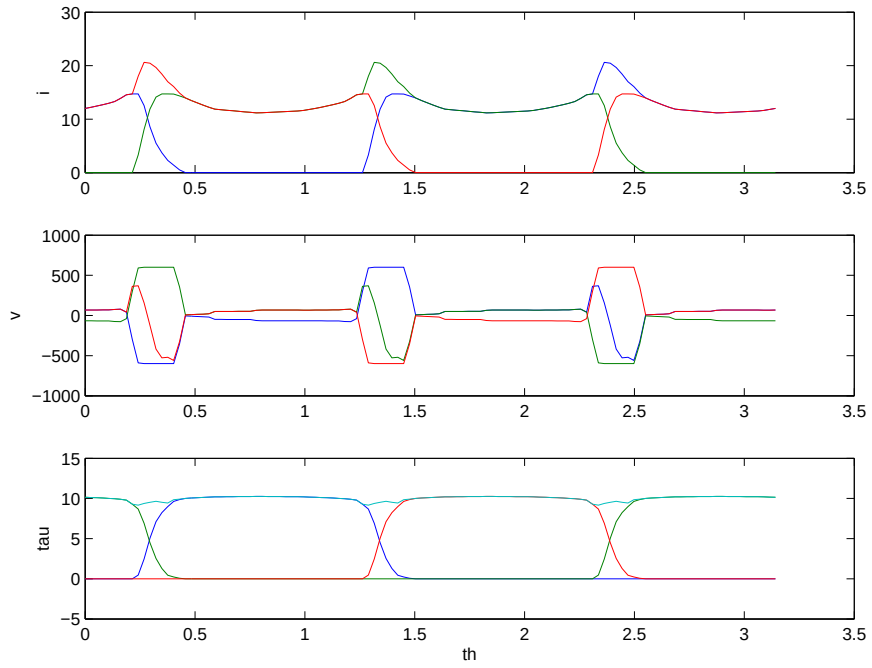

Figure 4: The optimal current, voltage, and torque waveforms for $\omega = 1000$ rpm.

Note that while the voltage limit is active over the commutation period for two of the windings, the third winding current is manipulated to maintain near-constant torque. Due to the fully pitched winding pattern, and in contrast to a motor with concentrated windings, it is not possible to derive a "torque-sharing function" for this example. This is because torque cannot be decomposed into components attributable to each winding (indeed, it is the changing mutual inductance between windings which generates torque; for a discussion of this, see [Mec93]).

High-speed operation. Here we show the optimal symmetric current, voltage, and phase torque waveforms for $\omega = 4000$ rpm, with all other values kept the same as for the low-speed example.

We see that when the rotor speed is higher, the optimal waveforms are much more complicated than the corresponding low-speed waveforms. Indeed, for high rotor speeds, deriving a simple, "closed-form" solution for the optimal current waveforms seems unlikely. The optimal current waveforms are strictly positive during the entire cycle; *i.e.*, there is no "firing angle" at which a given phase phase should be energized.

6.3 Computational aspects

To compute the waveforms above, problem (8) was solved using CVX [GB14, GB08], with Gurobi as the backend solver. We used a Linux machine with an Intel Xeon processor. The (global) solve time for low-speed operation was 6 seconds, while the solve time for high-speed operation was 4 minutes.

In order to verify that these solve times were representitive of the solve time over several motor operating points, we solved (8), using the same motor parameters (except for a courser grid of

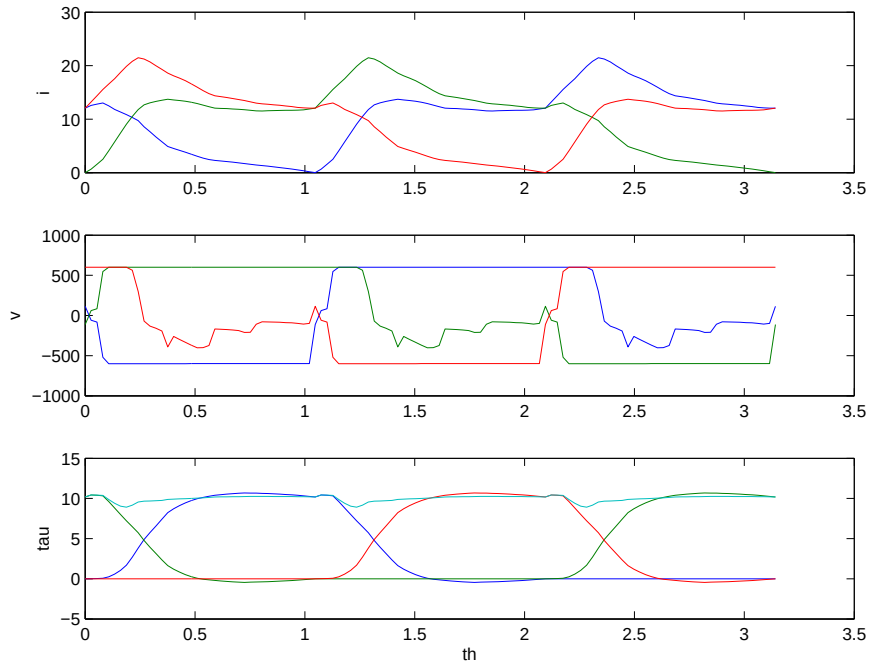

Figure 5: The optimal current, voltage, and torque waveforms for $\omega = 4000$ rpm.

$T = 20$) for 250 randomly selected values of ω and τ_{des}. in the intervals $[0, 4000]$ rpm and $[0, 15]$ N, respectively. The median global solve times was ten seconds, though we note that for some areas of operation (in particular, for very-low-torque operation) the global solve times can be very high, sometimes exceeding thirty minutes; in these cases, it is standard to terminate the solver early (after, say, several minutes) simply use the best point found so far.

7 Conclusion

In this paper we presented a method for generating optimal current waveforms for switched reluctance motors. Our model handles a reasonably complicated magnetic structure, and respects voltage and current constraints. Our method finds, to within reasonable accuracy, the globally optimal waveforms. Although we provide to runtime guarentees, for our example, optimal waveforms could be generated in a few minutes, raising the possibility that optimal waveforms can be generated and stored as a lookup table indexed by desired torque and rotor speed.

References

[Bal79] E. Balas. Disjunctive programming. *Annals of Discrete Mathematics*, 5:3–51, 1979.

[BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[CKKP02] C. Choi, S. Kim, Y. Kim, and K. Park. A new torque control method of a switched reluctance motor using a torque-sharing function. *IEEE Transactions on Magnetics*, 38(5):3288–3290, 2002.

[CMH93] J. Y. Le Chenadec, B. Multon, and S. Hassine. Current feeding of switched reluctance motor. optimization of the current waveform to minimize the torque ripple. *Proceedings of the 4th International Conference on Computational Aspects of Electromechanical Energy Converters and Drives*, 1993.

[CS99] S. Ceria and J. Soares. Convex programming for disjunctive convex optimization. *Mathematical Programming*, 86(3):595–614, 1999.

[CS02] P. L. Chapman and S. D. Sudhoff. Design and precise realization of optimized current waveforms for an 8/6 switched reluctance drive. *IEEE Transactions of Power Electronics*, 17(1):76–83, 2002.

[DK84] C. A. Desoer and E. S. Kuh. *Basic circuit theory*. Tata McGraw-Hill Education, 1984.

[FG07] A. Frangioni and C. Gentile. SDP diagonalizations and perspective cuts for a class of nonseparable MIQP. *Operations Research Letters*, 35(2):181–185, 2007.

[GB08] Michael Grant and Stephen Boyd. Graph implementations for nonsmooth convex programs. In V. Blondel, S. Boyd, and H. Kimura, editors, *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, pages 95–110. Springer-Verlag Limited, 2008.

[GB14] Michael Grant and Stephen Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. `http://cvxr.com/cvx`, March 2014.

[GL12] O. Günlük and J. Linderoth. Perspective reformulation and applications. In *Mixed Integer Nonlinear Programming*, pages 61–89. Springer, 2012.

[GT13] I. E. Grossmann and F. Trespalacios. Systematic modeling of discrete-continuous optimization models through generalized disjunctive programming. *American Institute of Chemical Engineers Journal*, 59(9):3276–3295, 2013.

[KSAS03] M. Kaiserseder, J. Schmid, W. Amrhein, and V. Scheef. Current shapes leading to positive effects on acoustic noise of switched reluctance drives. *International Journal for Computation and Mathematics in Electrical and Electronic Engineering*, 22(4):998–1008, 2003.

[KT06] J. Kleinberg and E. Tardos. *Algorithm design*. Pearson Education, 2006.

[LS94] H. C. Lovatt and J. M. Stephenson. Computer-optimised current waveforms for switched-reluctance motors. *IEE Proceedings Electric Power Applications*, 141(2):45–51, 1994.

[LS97a] H. C. Lovatt and J. M. Stephenson. Computer-optimised smooth-torque current waveforms for switched-reluctance motors. *IEE Proceedings Electric Power Applications*, 144(5):310–316, 1997.

[LS97b] H. C. Lovatt and J. M. Stephenson. Optimum excitation of switched reluctance motors. In *Eighth International Conference on Electrical Machines and Drives*, pages 356–360. IET, 1997.

[MB15] N. Moehle and S. Boyd. Optimal current waveforms for brushless permanent magnet motors. *International Journal of Control*, 88(7):1389–1399, 2015.

[Mec93] B. C. Mecrow. Fully pitched-winding switched-reluctance and stepping-motor arrangements. *IEEE Proceedings Electric Power Applications*, 140(1):61–70, 1993.

[MK03] C. Mademlis and I. Kioskeridis. Performance optimization in switched reluctance motor drives with online commutation angle control. *IEEE Transactions on Energy Conversion*, 18(3):448–457, 2003.

[MWC01] B. C. Mecrow, C. Weiner, and A. C. Clothier. The modeling of switched reluctance machines with magnetically coupled windings. *IEEE Transactions on Industry Applications*, 37(6):1675–1683, 2001.

[NW06] J. Nocedal and S. Wright. *Numerical optimization*. Springer Science & Business Media, 2006.

[TK12] H. A. Toliyat and G. B. Kliman. *Handbook of electric motors*, volume 120. CRC press, 2012.

[TMBB15] R. Takapoui, N. Moehle, S. Boyd, and A. Bemporad. A simple effective heuristic for embedded mixed-integer quadratic programming. *arXiv preprint arXiv:1509.08416*, 2015.

[VZP+07] M. Vasak, D. Zarko, N. Peric, F. Kolonic, and C. Hao. Bounding the torque ripple in switched reluctance motors using polyhedral invariant set theory. In *46th IEEE Conference on Decision and Control*, pages 6106–6111. IEEE, 2007.

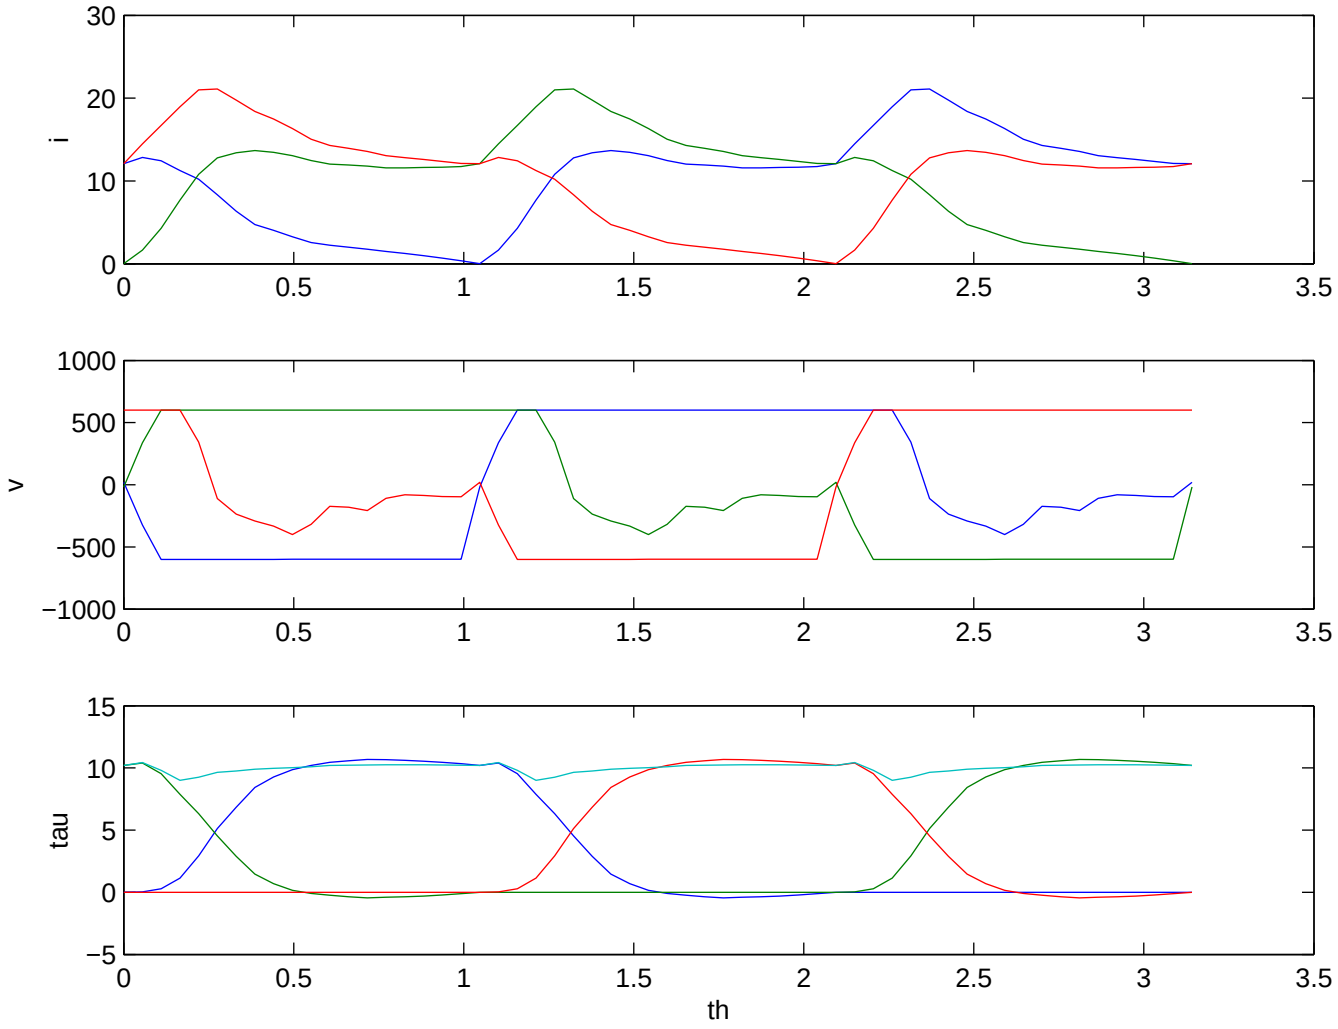
i
v
tau
th
30
20
10
0
1000
500
−500
−1000
15
5
−5
0
0.5
1
1.5
2
2.5
3
3.5